

REMARKS/ARGUMENTS

Applicants wish to thank the Examiner for her time to discuss this application with Applicants' attorney. These amendments and remarks are in response to the discussion and to the Final Office Action dated February 3, 2005. In the event the Examiner is not persuaded by Applicants' arguments, Applicants respectfully request that the Examiner enter the amendments to clarify issues upon appeal.

Claims 1-21 are pending in the present application. Applicants have added new claims 22, 23 and 24. Accordingly, claims 1-24 are pending.

New Claims

Applicants have added claims 22, 23 and 24. Claim 22 incorporates, in their entirety, claims 1, 2 and 3. Claim 23 incorporates claims 9, 10 and 11, and claim 24 is a system claim be consistent with claims 22 and 23. Because the new claims merely incorporate **existing claims**, Applicants respectfully submit that the amendments can be entered at this stage, and that they do not present new matter requiring a new search.

Claim Rejections

In the Final Office Action, the Examiner rejected claims 1-21 under 35 U.S.C. §103(a) as being unpatentable over Josten et al. (U.S. Patent No. 5,574,902) in view of Ponnekanti (U.S. Patent No. 6,591,269). In so doing, the Examiner stated:

Regarding claim 1, Josten teaches a method for optimizing command execution in a database system, wherein data records are stored on a plurality of data pages therein (col. 5, line 65-col. 6, line 6), the method comprising the steps of:

(a) 'providing an identifier to each data page' as an ordinal number (ORD#) is assigned to a data page buffer control block (BCB) in dirty page list (DPL) (col. 7, lines 10-16 and 42-52; col. 11, lines 44-46, and Fig. 2, element 44 (i.e., ordinal #));

(b) 'selecting a data record from a data page' as list 42 includes a series of the ordinal number from DPL that are associated with data pages in the LCB that were accessed by the transaction corresponding to TPL-1 (col. 7, line

53 – col. 8, line 2; col. 5, lines 5-10; col. 7, lines 10-16; col. 8, lines 1-7; col. 8, lines 46-51);

(c) ‘copying the selected data record to a second storage area’ as the data manager issues a SETWRITE request to indicate intent to update the named data page (col. 7, lines 10-16; col. 8, lines 9-18);

(d) ‘verifying that the selected data record has not been modified since the time that it was copied to the second storage area based upon the identifier’ as a test is made to detect a consecutive update to the same data page by comparing the ORD# of the last entry of the transaction page list (TPL) with the ORD# of the new entry in the buffer control block (col. 11, lines 28-38); and

(e) ‘executing the command’ as committing transactions schedules write I/Os for all TPL entries (col. 18, lines 3-7).

a). Josten does not explicitly teach the identifier indicating when any of the data records contained therein were last modified.

Ponnekanti, however, teaches the log records contain only the PAGEIDs and the timestamps of the source page and the target page and the positions of the first and last key that were copied (col. 11, lines 47-49).

It would have been obvious to one of ordinary skill in the art at the time of the invention was made to combine the teachings of the cited references because Ponnekanti’s teaching would have allowed Josten’s to enhance the speed in which the database server stores, retrieves, and processes particular data records as indicated by Ponnekanti at col. 3, lines 1-10 and col. 19, lines 40-50.

Regarding claims 2 and 10, Josten further teaches wherein the copying step (c) includes:

(c1) copying and storing the identifier to the second storage area (col. 4 lines 27-32 and col. 7, line 57-col. 8, line 2).

Regarding claims 3 and 11, Josten further teaches wherein the verifying step (d) includes:

(d1) determining a current identifier for the data page (col. 11, lines 28-31);

(d2) comparing the current identifier with the stored identifier (col. 11, lines 34-36); and

(d3) concluding the selected data record has not been modified when the current identifier is the same as the stored identifier (col. 11, lines 36-38).

Applicants respectfully disagree.

The present invention is directed to optimizing command execution in a relational database management system RDBMS where data records are stored in data pages. According to the present invention, *each* data page, not just dirty data pages, is provided with an identifier which indicates when the page was modified last. Thus, whenever a record on the data page is

modified, the identifier necessarily changes. When the RDBMS receives a command to access a data record on a data page, such as an UPDATE or DELETE command, the system copies the selected *data record* from the page, along with the page's identifier, and stores them temporarily in a secondary storage area.

The system uses the stored identifier to verify that the copied record has not been modified during the period in which it was copied and stored. For instance, the system compares the stored identifier with a current identifier for the page. If the page has *not been modified*, the stored identifier and the current identifier will be the same, and the system can conclude that the stored data record matches the current record in the database. Accordingly, when the stored identifier is the same as the current identifier, the system is not required to compare the data record in the temporary copy with the record in the current table, which can be a costly event if the records are extensive.

By providing *each data page* with the identifier and then storing the identifier with the data record in the secondary storage, the RDBMS can quickly and easily determine whether the data page containing the stored data record was modified by comparing identifiers. Accordingly, the present invention minimizes the number of times the RDBMS must go into the data page itself to verify that the stored record is consistent with the current record.

In contrast, Josten is directed to a method for destaging updated locally cache pages in a multisystem shared disk environment, such as a Sysplex, with a high-speed shared electronic store. In Josten, each computer processing complex (CPC) caches data pages in its local cache buffer (LCB). When one of those pages is updated, it eventually must be written to shared stable storage ("externalized") and cross-invalidated in other CPCs when a transaction commits. (Col. 4, lines 8-13). Josten presents (a) an efficient process for identifying the set of cached pages that must be externalized and (b) a fast process for scheduling of the externalizing write I/Os so that

the corresponding transaction can release its locks. (Col. 4, lines 13-18). To do this, Josten assigns an ordinal number (ORD#) to a *dirty data page*, which is defined as a page that contains a record that has been updated by a transaction, but has not yet been written to external stable storage. Josten then places the ORD# in a list associated with the transaction (transaction page list (TPL)). (Col. 7, lines 1-18).

Ponnekanti is directed to performing an online rebuild of a B+ tree index by copying the index rows to newly allocated pages in the key order so that good space utilization and clustering are achieved. (Abstract). In Ponnekanti, index rows are copied to newly allocated pages in the key order, and old pages are deallocated. Multiple leaf pages are rebuilt and then changes to higher levels are propagated. While propagating leaf level changes to higher levels, level 1 pages are reorganized. (Abstract). Thus, the need for a separate pass is eliminated.

Claims 1, 9 and 17 are Allowable

Applicants respectfully submit that Josten and Ponnekanti fail to teach or suggest “providing an identifier to *each* data page” in a database system, as recited in claims 1 and 9, and “means for providing an identifier on *each* data page” in the database system, as recited in claim 17, where the identifier indicates “when any of the data records contained therein were last modified.”

According to the Final Office Action, the Examiner suggests that the identifier of the present invention is taught by Josten’s ORD# which is assigned to a *dirty* data page in the local cache buffer (LCB). (Col. 7, lines 10-16). Thus, a clean data page is not assigned an ORD#. In the present invention, the identifier is provided to *each data page*, regardless of whether the page is clean or dirty. Accordingly, Applicants respectfully submit that Josten’s ORD# is not equivalent to the identifier of the present invention.

Moreover, the identifier indicates *when* any of the data records contained in the data page was modified last. The Examiner concedes that the ORD# does not provide such an indication, but suggests that Ponnekanti's log records do at column 11, lines 47-49. Applicants respectfully disagree. In Ponnekanti, the application runs as a sequence of transactions, with each transaction rebuilding leaf pages in the page chain. At the end of each transaction, the new pages generated are flushed to disk and then the old pages that were removed from the tree are made available for fresh allocations. The log records contain PAGEIDs and timestamps of the source page and the target page, as well as positions of a first and a last key copied. (Column 11, lines 30-49).

The PAGEID and timestamp of a source or target page has no relation to Josten's ORD#. In addition, nothing in Ponnekanti teaches or suggests that the timestamp of the source or target page indicates "when any of the data records contained therein were last modified," as recited in claims 1, 9 and 17. Accordingly, Applicants respectfully submit that the combination of Josten and Ponnekanti fail to teach or suggest the identifier of the present invention, as recited in claims 1, 9 and 17.

In addition, Applicants respectfully maintain that neither reference teaches or suggests "verifying that the selected data record has not been modified since the time that it was copied to the second storage area *based upon the identifier*," as recited in claims 1 and 9, and "means for verifying that the selected data record has not been modified since the time that it was copied to the second storage area by determining that the stored identifier is the same as the current identifier from the data page," as recited in claim 17. In Josten, the ORD# for a dirty page does not change once the page becomes dirty. If a transaction updates a data page *more than once*, the ORD# for the dirty data page does not change. Thus, in Josten, it is not possible to verify that "the selected data record has not been modified since the time that it was copied to the second storage area *based upon the identifier*," as recited in claims 1 and 9, and Josten does not teach or

suggest “means for verifying that the selected data record has not been modified since the time that it was copied to the second storage area by determining that the stored identifier is the same as the current identifier from the data page,” as recited in claim 17.

In the Final Office Action, the Examiner states that this feature is taught by Josten at column 11, lines 28-38. Applicants respectfully disagree. In the cited portion, Josten discusses avoiding duplicate TPL entries by determining whether a consecutive update has been performed on the same data page by comparing the ORD# of the last entry of the TPL with the ORD# of the new entry. If the ORD#s are *different*, the new entry is made at the end of the TPL. This process does not verify “that the selected data record has not been modified since the time that it was copied to the second storage area,” as recited in claims 1, 9 and 17.

For the foregoing reasons, Applicants respectfully submit that claims 1, 9 and 17 are allowable over Josten and Ponnekanti. Claims 2-8, 10-16, and 18-21 depend on claims 1, 9 and 17, respectively, and therefore, the above arguments apply with full force. Thus, claims 2-8, 10-16, and 18-21 are also allowable over Josten in view of Ponnekanti.

Claims 22, 23 and 24 are Allowable

As stated above, claim 22 incorporates claims 1, 2 and 3, claim 23 incorporates claims 9, 10 and 11, and claim 24 is a system claim be consistent with claims 22 and 23. As such, claims 22, 23 and 24 are allowable over Josten and Ponnekanti for the reasons set forth above. In addition, claims 22-24 are allowable because neither of the references teaches or suggests steps (d1), (d2) and (d3). In the Final Office Action, the Examiner asserts that Josten teaches steps (d1)-(d3) at column 11, lines 28-38. The cited portion states,

Efficient tracking of data pages in a data base that are modified by a transaction is provided by the transaction page list (TPL) of this invention, each entry of which consists of the associated ORD # only. Before adding an entry to the TPL for a transaction, a simple test is made to detect a consecutive update to the same data page by comparing the ORD # of the last entry of the TPL with the ORD # of the new entry, which is available in the buffer-control block (BCB) for

the data page copy in LCB. If the last TPL entry is not the same as the entry for ORD #, the new entry is made at the end of the TPL. Consecutive duplicate TPL entries are thereby avoided, which reduces the TPL size without significant TPL search activity.

In Josten, if a clean data page becomes dirty (i.e., it's modified by a first transaction), an ORD# is assigned to the dirty page, and the ORD# is stored in the dirty page list (DPL).

According to Josten, "[t]he ORD# associated with a data page in LCB remains the same until the data page is written to stable storage." Column 11, lines 24-25. Thus, once a clean page becomes dirty, it remains dirty until it is externalized. Subsequent modifications to the same dirty page do not change the dirty page's ORD#.

The portion cited above describes how the size of the TPL can be reduced by eliminating consecutive duplicate TPL entries. As stated in Josten, each TPL entry "consists of the associated ORD# only." (Column 11, line 31). If two consecutive transactions operate on the same data page, the ORD# for the first transaction is the same ORD# as the second transaction. Without Josten's "test," the TPL would have two consecutive entries that are duplicates. With the "test," consecutive duplicate TPL entries are avoided.

In the present invention, the current identifier for the data page is compared to the stored identifier for the data page. If the current identifier is the **same** as the stored identifier, the present invention concludes that the data record **has not been modified**. In the present invention, this conclusion is logical because the identifier for the data page changes if one of the data records in the data page has been modified. If the identifier for the data page is the same, none of the records has been modified.

In contrast, Josten's ORD# for a data page "remains the same" regardless of the number of times the data page has been modified. If two consecutive transactions modify the same page, the ORD#'s for the first and second transaction are the same. Thus, if the second ORD# (current identifier) is the same as the first ORD# (stored identifier), Josten concludes that the second

transaction **has modified** the same data page. This is in direct contrast to the present invention which teaches that the data record in a data page **has not been modified** when the identifiers are the same.

Based on a close comparison of Josten to the present invention, Applicants submit that claims 22, 23 and 24 are allowable over Josten in view of Ponnekanti for this additional and alternative reason.


Conclusion

In view of the foregoing, Applicants submit that claims 1-24 are allowable over the cited references. Applicants respectfully request reconsideration and allowance of the claims as now presented.

Applicants' attorney believes that this application is in condition for allowance. Should any unresolved issues remain, Examiner is invited to call Applicant's attorney at the telephone number indicated below.

Respectfully submitted,

SAWYER LAW GROUP LLP



Stephen G. Sullivan
Attorney for Applicant(s)
Reg. No. 38,329
(650) 493-4540

April 15, 2005

Date